

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 1 353 264 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
15.10.2003 Bulletin 2003/42

(51) Int Cl.7: G06F 3/06

(21) Application number: 02017669.9


(22) Date of filing: 06.08.2002

<div>(84) Designated Contracting States: AT BE BG CH CY CZ DE DK EE ES FI FR GB GR IE IT LI LU MC NL PT SE SK TR Designated Extension States: AL LT LV MK RO SI</div> <div>(30) Priority: 09.04.2002 JP 2002106262</div> <div>(71) Applicant: Hitachi, Ltd. Chiyoda-ku, Tokyo 101-8010 (JP)</div>	<div>(72) Inventors:<ul style="list-style-type: none">• Tanaka, Katsuya, Hitachi, Ltd., Int. Prop. Group Chiyoda-ku, Tokyo 100-8220 (JP)• Fujimoto, Kazuhisa, Hitachi, Ltd., Int Prop. Group Chiyoda-ku, Tokyo 100-8220 (JP)</div> <div>(74) Representative: Strehl Schübel-Hopf & Partner Maximilianstrasse 54 80538 München (DE)</div>
-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

(54) Disk storage system having disk arrays connected with disk adaptors through switches

(57) A disk storage system has high throughput between a disk adapter (DKA) of a disk controller (DKC) and a disk array (DA). The disk adapter (DKA) of the disk controller (DKC) is connected to the disk array (DA) through switches (SW1, SW2, SW3, SW4). Data on a channel (D11, D12, D13, D14) between the switch (SW1) and a RAID group (R1) is multiplexed in the switch (SW1) to be transferred onto a channel (D01) between the switch (SW1) and the disk adapter (DKA) and data on the channel (D01) between the switch (SW1) and the disk adapter (DKA) is demultiplexed in the switch (SW1) to be transferred onto the channel (D11, D12, D13, D14) between the switch (SW1) and the RAID group (R1). A data transfer rate on the channel (D01) is made higher than that on the channel (D11, D12, D13, D14).

FIG.1



EP 1 353 264 A2

Description

BACKGROUND OF THE INVENTION

[0001] The present invention relates to a secondary storage apparatus in a computer system and more particularly to a disk storage system having high input/output data transfer performance.

[0002] In a current computer system, data required by a CPU (Central Processing Unit) are stored in a secondary storage apparatus and data are written in and read out from the secondary storage apparatus when the CPU or the like requires the data.

[0003] As the secondary storage apparatus, a nonvolatile memory medium is generally used and as its representative, there is a disk storage system such as a magnetic disk apparatus and an optical disk.

[0004] With the high advancement of the information-oriented society in recent years, high-performance of the secondary storage apparatus of this kind is required in the computer system.

[0005] Fig. 9 is a block diagram schematically illustrating a conventional disk storage system.

[0006] In Fig. 9, the disk storage system includes a disk controller DKC and a disk array DA.

[0007] The disk controller DKC includes a channel adapter CHA for connecting the disk storage system to a higher-rank CPU (not shown), a cache memory CM for temporarily storing data written in and read out from the disk array DA and a disk adapter DKA for connecting the disk controller DKC to the disk array DA.

[0008] The channel adapter CHA, the cache memory CM and the disk adapter DKA are connected through a bus or switch.

[0009] The channel adapter CHA includes four channels C1, C2, C3 and C4 connected to the CPU.

[0010] The disk adapter DKA is connected to the disk array through four channels D1, D2, D3 and D4.

[0011] The disk array DA includes disk groups R1, R2, R3 and R4.

[0012] When the RAID system is constructed in the disk array DA, the disk groups R1, R2, R3 and R4 each constitute RAID groups.

[0013] Write data inputted from the channels C1, C2, C3 and C4 are stored in the cache memory CM and at the same time the write data are divided into data of block-size unit, so that the data divided in the block unit are sent through three channels of the channels D1, D2, D3 and D4 and a parity calculated from the divided data is sent through a remaining channel from the disk adapter DKA to the disk array DA.

[0014] When data is read out, it is first examined whether the data is present or stored in the cache memory CM.

[0015] When the data is present in the cache memory CM, the data is read out through the channel adapter CHA from the cache memory CM and transmitted through the channel adapter CHA to the CPU.

[0016] When the data is not present in the cache memory CM, the disk adapter DKA reads out the data divided in the block unit from the disk array DA through the channels D1, D2, D3 and D4 and transmits the read data through the channel adapter CHA to the CPU. The conventional technique of this kind is named the first conventional technique.

[0017] The disk storage system related to the first conventional technique is described in, for example, "MAIN FRAME in Separate Volume of Nikkei Computer '98", pp. 144 to 153 issued by Nikkei BP Co. (1998).

[0018] A disk storage system having a disk array connected to a disk adapter through switches is disclosed in JP-A-5-173722 entitled "Exchange Device of Multi-Channel Data and Parity".

[0019] The conventional technique described in the above publication is hereinafter named the second conventional technique.

[0020] According to the second conventional technique, the number of buses related to the disk array and the number of buses related to the disk adapter can be set up independently.

[0021] A disk storage system having a disk array connected to a disk adapter through a buffer control block is disclosed in JP-A-6-19627 entitled "Rotation Type Storage Apparatus".

[0022] The conventional technique described in the above publication is hereinafter named the third conventional technique.

[0023] According to the third conventional technique, a data transfer rate between the disk adapter and the disk array can be set up to any value and influence due to a waiting time caused by rotation of the disk can be reduced.

[0024] The data transfer rate per channel is increased year after year with the progress of network technique.

[0025] For example, the data transfer rate per channel in a fiber channel used in a disk storage system is 1 to 2 Gbps at present but is planned to be increased to 4 to 10 Gbps in the near future.

[0026] It is forecasted that the throughput between the CPU and the channel adapter (hereinafter named the front-end) complies with the above increased data transfer rate.

[0027] However, it is forecasted that the throughput between the disk adapter and the disk array (hereinafter named the back-end) is not increased as the front-end due to the following reasons.

[0028] The first reason is that the disk drive contains mechanical components and accordingly the high-speed operation is difficult as compared with the front-end including only electronic and optical elements as elements influencing the high-speed operation.

[0029] The second reason is that the cost of the disk storage system having a large number of disk drives is increased when high-speed interfaces are mounted for all of the disk drives even if operation of the disk drives is made fast.

[0030] The first conventional technique has a problem that the performance of the disk storage system is not improved due to detachment of the throughputs of the front-end and the back-end even if the data transfer rate per channel of the channel adapter is improved.

[0031] Further, it is considered that a large number of low-speed ports are provided in the disk adapter in order to improve the throughput of the back-end, although the increased ports in the disk adapter complicates control.

[0032] In the second conventional technique, switches can be provided between the disk adapter and the disk array to thereby increase the number of ports for increased disks, although there is a problem that the data transfer rate per channel is limited to that of the disk array and accordingly the throughput between the disk adapter and the disk array becomes the neck of performance.

[0033] The third conventional technique is a technique capable of reducing influence due to the waiting time caused by rotation of the disk but has a problem that detachment between the front-end and the back-end cannot be reduced.

SUMMARY OF THE INVENTION

[0034] It is an object of the present invention to provide a disk storage system having high throughput between a disk adapter and a disk array.

[0035] It is another object of the present invention to provide a disk storage system having high throughput between a disk adapter and a disk array and the increased number of disk drives connected thereto.

[0036] It is still another object of the present invention to provide a disk storage system having a highly reliable disk array.

[0037] It is still another object of the present invention to provide a disk storage system having a highly reliable network between a disk adapter and a disk array.

[0038] It is still another object of the present invention to provide a disk storage system having a network with high reliability and high throughput between a disk adapter and a disk array.

[0039] It is still another object of the present invention to provide a disk storage system capable of performing reading operation from and writing operation in a disk with high throughput.

[0040] It is still another object of the present invention to provide a disk storage system capable of keeping high throughput.

[0041] It is still another object of the present invention to provide an inexpensive disk storage system having high throughput.

[0042] In order to achieve the above objects, the disk storage system according to the present invention includes a disk controller and a disk array and the disk controller includes a channel adapter, a cache memory and a disk adapter.

[0043] The disk adapter is connected to the disk array

through a switch having a buffer memory and the data transfer rate per channel between the disk adapter and the switch is set to be higher than that between the switch and the disk array.

[0044] The switch changes over connection between ports connected to which the disk adapter is connected and ports to which disk drives constituting the disk array are connected in accordance with destination information in a frame for each of inputted frames.

[0045] Further, the disk array includes a plurality of disk drives connected into a loop and the disk adapter is connected to the plurality of disk arrays through switches having buffer memories. The data transfer rate per channel between the disk adapter and the switch is set to be higher than that between the switch and the plurality of disk arrays and the switch changes over connection between ports to which the disk adapter is connected and ports to which the plurality of disk arrays are connected in accordance with destination information in a frame for each of inputted frames.

[0046] Moreover, the disk adapter is connected to the disk array through switches having buffer memories and a combination of disk drives connected to the same switch constitutes a RAID group. The data transfer rate per channel between the disk adapter and the switch is set to be higher than that between the switch and the disk array and the switch changes over connection between ports to which the disk adapter is connected and ports to which the disk drives constituting the RAID group are connected in accordance with destination information in a frame for each of inputted frames.

[0047] Furthermore, the disk storage system includes first and second disk controllers and a plurality of disk arrays. The first disk controller includes a first channel adapter, a first cache memory and a first disk adapter and the second disk controller includes a second channel adapter, a second cache memory and a second disk adapter.

[0048] The first disk adapter is connected to the plurality of disk arrays through a first switch having a buffer memory and the second disk adapter is connected to the plurality of disk arrays through a second switch having a buffer memory. Further, the first switch is connected to the second disk adapter and the second switch is connected to the first disk adapter.

[0049] The data transfer rate per channel between the second disk adapter and the second switch and between the first disk adapter and the second switch is set to be higher than that between the second switch and the plurality of disk arrays.

[0050] The first switch changes over connection between ports to which the first or second disk adapter is connected and ports to which the plurality of disk arrays are connected in accordance with destination information in a frame for each of inputted frames and the second switch changes over connection between ports to which the first or second disk adapter is connected and ports to which the plurality of disk arrays are connected

5

EP 1 353 264 A2

6

in accordance with destination information in a frame for each of inputted frames.

[0051] Moreover, the first switch is connected to the second switch through a channel having the same data transfer rate as the channel connected between the first disk adapter and the second switch and a channel having the same data transfer rate as the channel connected between the second disk adapter and the first switch.

[0052] Further, when data is read out from the disk array, data transferred from the disk array to the switch is multiplexed in the switch to be transferred to the disk adapter and when data is written in the disk array, data transferred from the disk adapter to the switch is demultiplexed in the switch to be transferred to the disk array.

[0053] Furthermore, when data is written from the disk adapter into the disk array, the disk adapter sets destination information in a frame to be transmitted so that connection between the ports is changed over and when data is read out from the disk array to the disk adapter, the switch changes over connection between ports by the round robin method.

[0054] Further, the number of ports to be changed over is set to be substantially equal to the ratio of the data transfer rate per channel between the disk adapter and the switch to that between the switch and the disk array.

[0055] In addition, the disk adapter is connected to the switch through optical fiber cable and the switch is connected to the disk array through metal cable.

BRIEF DESCRIPTION OF THE DRAWINGS

[0056]

Fig. 1 schematically illustrates a disk storage system according to a first embodiment of the present invention;

Fig. 2 illustrates the configuration of a switch used in the present invention;

Fig. 3 illustrates the configuration of a switch used in the present invention;

Fig. 4 is a diagram illustrating operation of the switch used in the present invention;

Fig. 5 is a diagram illustrating operation of the switch used in the present invention;

Fig. 6 is a diagram illustrating a method of increasing disk drives in the first embodiment of the present invention;

Fig. 7 schematically illustrates a disk storage system according to a second embodiment of the present invention;

Fig. 8 schematically illustrates a disk storage system according to a third embodiment of the present invention;

Fig. 9 schematically illustrates a conventional disk storage system;

Fig. 10 is a diagram explaining a connection form by FC-AL;

Fig. 11 is a diagram explaining a connection form by FC-AL; and

Fig. 12 schematically illustrates a disk storage system according to a fourth embodiment of the present invention.

DESCRIPTION OF THE EMBODIMENTS

[0057] Embodiments of the present invention are now described in detail with reference to the accompanying drawings.

[0058] Fig. 1 schematically illustrates a disk storage system according to a first embodiment of the present invention.

[0059] The disk storage system of the embodiment includes a disk controller DKC and a disk array DA.

[0060] The disk controller DKC includes a channel adapter CHA, a cache memory CM and a disk adapter DKA.

[0061] The channel adapter CHA performs control required when data is transmitted and received between a higher-rank CPU (not shown) and the disk controller DKC.

[0062] C1, C2, C3 and C4 represent channels used when the channel adapter CHA communicates with the CPU.

[0063] The cache memory CM is a memory for temporarily storing data inputted/outputted by the disk storage system of the embodiment.

[0064] The disk adapter DKA performs control required when data is transmitted and received between the disk controller DKC and the disk array DA.

[0065] The disk adapter DKA is connected to the disk array DA through channels D01, D02, D03 and D04.

[0066] The disk adapter DKA and the disk array DA can perform the full-duplex operation through the channels D01, D02, D03 and D04.

[0067] In the disk storage system of the embodiment, the disk adapter DKA and the disk array DA are connected through switches SW1, SW2, SW3 and SW4.

[0068] The disk array DA is composed of disk groups R1, R2, R3 and R4.

[0069] The disk group R1 is connected to the disk adapter DKA through the switch SW1.

[0070] Similarly, the disk groups R2, R3 and R4 are connected to the disk adapter DKA through the switches SW2, SW3 and SW4, respectively.

[0071] In the disk storage system of the embodiment, when the RAID system is constructed, the disk groups R1, R2, R3 and R4 are each the RAID group.

[0072] In the embodiment, the RAID group is constituted by four disk drives but the number of disk drives constituting the RAID group is not limited to four.

[0073] The flow of data in reading and writing of data to each disk group is now described by taking the disk group R1 as an example. R1 is the RAID group of RAID level 5.

[0074] Data transmitted from the CPU through the

channels C1, C2, C3 and C4 in order to be written in the disk group R1 is divided into data of block unit in the disk adapter DKA and at the same time a parity is produced from the data divided into the block unit.

[0075] The data divided into the block unit and the parity are supplied through the channel D01 to the switch SW1.

[0076] The switch SW1 performs routing of the data divided into the block unit and the parity in accordance with the RAID control to be distributed to the channels D11, D12, D13 and D14

[0077] In reading of data, the disk adapter DKA reads out data divided into the block unit from the disk group R1 through the channels D11, D12, D13 and D14 and converts the data into serial data in the switch SW1 to receive the read-out data through the channel D01.

[0078] In the conventional disk storage system shown in Fig. 9, data and a parity thereof to be written in the disk array are already distributed on separate channels D1, D2, D3 and D4 connected to the disk adapter DKA.

[0079] On the contrary, the disk storage system of the embodiment is different from the conventional system in that the data and the parity are distributed on separate channels after the passage of the switch SW1.

[0080] Operation of the switch characterized in the disk storage system of the embodiment is next described by taking the switch SW1 as an example. Operation of the switches SW2 to SW4 is also the same as that of the switch SW1.

[0081] As shown in Fig. 1, the switch SW1 includes input/output ports P1, P2, P3, P4 and P5.

[0082] The ports P1, P2, P3, P4 and P5 are input/output ports in which the full-duplex operation can be made and each include a buffer memory.

[0083] The internal configuration of the switch SW1 is shown in Figs. 2 and 3.

[0084] For simplification, description is made in accordance with switching operation divided in the flowing direction of data.

[0085] Further, data flowing on the channels D01, D11, D12, D13 and D14 are transmitted and received in frame unit and encoded in accordance with the 8B10B encoding scheme.

[0086] Fig. 2 illustrates the circuit configuration in which the frame in a block is inputted from the port P1 and outputted from the ports P2, P3, P4 and P5.

[0087] This operation corresponds to switching operation in writing of data in the disk array.

[0088] As shown in Fig. 2, the switch SW1 includes a crossbar switch XSW and a switch controller CTL.

[0089] The crossbar switch XSW is of 5×5 and includes input ports in1, in2, in3, in4 and in5 and output ports out1, out2, out3, out4 and out5.

[0090] The frame inputted from the port P1 is supplied through a serial-to-parallel converter SP1, a buffer memory BM1 and an 8B10B decoder DEC1 to the switch controller CTL and the input port in1.

[0091] The switch controller CTL decodes a destina-

tion address written in a header portion of the input frame and changes over the crossbar switch XSW.

[0092] By way of example, when the port P2 is selected as the output port, the inputted frame passes through the output port out2, an 8B10B encoder ENC2, a buffer memory BM2 and a parallel-to-serial converter PS2 and is outputted from the port P2.

[0093] The buffer memories BM1 and BM2 are FIFO (First-In First-Out) memories.

[0094] The serial-to-parallel converter SP1 converts the 8B10B-encoded serial data into parallel data of 10-bit width and writes the parallel data into the buffer memory BM1 in synchronism with the data transfer rate equal to one tenth of that at the port P1.

[0095] The 8B10B decoder DEC1 reads out the 10-bit parallel data from the buffer memory BM1 in synchronism with the operation speed of the crossbar switch XSW and subjects the 10-bit parallel data to 8B10B decoding to convert it into 8-bit parallel data.

[0096] The 8B10B encoder ENC2 encodes the 8-bit parallel data switched by the crossbar switch XSW into 8B10B-encoded data again to convert it into 10-bit parallel data and writes the data into the buffer memory BM2 in synchronism with the operation speed of the crossbar switch XSW.

[0097] The parallel-to-serial converter PS2 reads out the 10-bit parallel data from the buffer memory BM2 in synchronism with the data transfer rate equal to one tenth of that at the port P2 and converts the data into serial data to be outputted from the port P2.

[0098] Thus, the switch SW1 converts the data transfer rate at the port P1 into that at the port P2.

[0099] Fig. 4 is a diagram showing waveforms of the frame inputted to the port P1 and the frame outputted from the ports P2, P3, P4 and P5.

[0100] The high level of the waveforms represents the time that the frame is present and the low level thereof represents the time that the frame is not present.

[0101] The frame has the frame length changing in accordance with the capacity of data to be transmitted, while in this example the sequential access to the disk array is made and the frame length is fixed.

[0102] In Fig. 4, the data transfer rate at the input port P1 is m times of that at the output ports P2, P3, P4 and P5.

[0103] Accordingly, the duration T1 for the frame Fb2 at the port P1 is extended to the duration T3 upon outputted from the port P2.

[0104] Thus, $T3 = m \times T1$.

[0105] When the data transfer rate for input is high and the data transfer rate for output is low, the buffer memory at the output port is overflowed unless the switch is changed over periodically, so that the throughput is reduced.

[0106] In order that the frame passes through the switch without reduction of the throughput, it is necessary to change over the output port periodically as shown in Fig. 4.

9

EP 1 353 264 A2

10

[0107] When the number of ports to be changed over is n , the changing-over period of switch is $T2 \approx n \times T1$ (the time having no frame is neglected).

[0108] When $T2 \geq T3$, the frames do not collide and the throughput is not reduced.

[0109] $T2 \geq T3$ is the same as $n \geq m$.

[0110] In other words, the condition that the throughput is not reduced in the switch upon writing of data to the disk array is to set the number n of ports of the switch to be changed over periodically to be equal to or larger than a ratio m of the data transfer rate per channel between the disk adapter and the switch to that between the switch and the disk array.

[0111] When this condition is satisfied, data inputted from the port P1 is subjected to rate conversion in the buffer memory and changed over in frame unit periodically to thereby be demultiplexed by means of the switch, so that the demultiplexed data are distributed into the ports P2, P3, P4 and P5 to be outputted.

[0112] One of methods of changing over the switch periodically is to make the disk group connected to the switch to be the RAID group.

[0113] The switch is changed over periodically in accordance with the striping control of the RAID.

[0114] Fig. 3 schematically illustrates the circuit configuration in which the frame is inputted from the ports P2, P3, P4 and P5 and outputted from the port P1.

[0115] This operation corresponds to switching operation in reading of data from the disk array.

[0116] For example, the frame inputted from the port P2 is supplied through a serial-to-parallel converter SP2, a buffer memory BM2 and an 8B10B decoder DEC2 to the switch controller CTL and the input port in2.

[0117] The switch controller CTL decodes a destination address written in a header portion of the input frame and changes over the crossbar switch XSW.

[0118] In the case of Fig. 3, data inputted from the ports P2, P3, P4 and P5 successively while changing over the crossbar switch XSW in accordance with the round robin method are all outputted to the port P1. That is, the frames are supplied to the plurality of input ports (P2, P3, P4 and P5) simultaneously in the reading operation. The plurality of input frames are not required to be supplied to the input ports in synchronism with one another. The switch changes over connection between input and output ports in the round robin manner to thereby transfer the plurality of input frames to the output port (P1) frame by frame. With the round robin method, the switch is changed over periodically. Further, even in the reading operation, the switch must be changed over in accordance with destination information in the frame.

[0119] The frame passes through the output port out1, an 8B10B encoder ENC1, the buffer memory BM1 and a parallel-to-serial converter PS1 and is outputted from the port P1.

[0120] The serial-to-parallel converter SP2 converts the 8B10B-encoded serial data into parallel data of 10-bit width and writes it in the buffer memory BM2 in

synchronism with the data transfer rate equal to one tenth of that at the port P2.

[0121] The 8B10B decoder DEC2 reads out the 10-bit parallel data from the buffer memory BM2 in synchronism with the operation speed of the crossbar switch XSW and subjects the 10-bit parallel data to 8B10B decoding to convert it into 8-bit parallel data.

[0122] The 8B10B encoder ENC1 encodes the 8-bit parallel data switched by the crossbar XSW into 8B10B-encoded data again to convert it into 10-bit parallel data and writes it into the buffer memory BM1 in synchronism with the operation speed of the crossbar switch XSW.

[0123] The parallel-to-serial converter PS1 reads out the 10-bit parallel data from the buffer memory BM1 in synchronism with the data transfer rate equal to one tenth of that at the port P1 and converts the 10-bit parallel data into serial data to be outputted from the port P1.

[0124] Thus, the switch SW1 converts the data transfer rate at the port P2 into that at the port P1.

[0125] Fig. 5 is a diagram showing waveforms of the frame inputted to the ports P2, P3, P4 and P5 and the frame outputted from the port P1.

[0126] The high level of the waveforms represents the time that the frame is present and the low level thereof represents the time that the frame is not present.

[0127] The frame has the frame length changing in accordance with the capacity of data to be transmitted, while in this example the sequential access to the disk array is performed and the frame length is fixed.

[0128] In Fig. 5, the data transfer rate at the input port P1 is m times of that at the output ports P2, P3, P4 and P5.

[0129] Accordingly, the duration $T4$ for the frame Fe5 at the port P5 is contracted to the duration $T5$ upon outputted from the port P1.

[0130] Thus, $T4 = m \times T5$.

[0131] The duration required to output the frames Fe2, Fe3, Fe4 and Fe5 from the port P1 is assumed to be $T6$.

[0132] When the number of ports to be changed over by the switch is n , $T6 \approx n \times T5$ (the time having no frame is neglected).

[0133] In order to prevent reduction of the throughput due to congestion in the switch, $T6 \leq T4$ is required. $T6 \leq T4$ is the same as $n \leq m$.

[0134] In other words, the condition that the throughput is not reduced in the switch in reading of data to the disk array is to set the number n of ports to be changed over periodically to be equal to or larger than a ratio m of the data transfer rate per channel between the disk adapter and the switch to the data transfer rate per channel between the switch and the disk array.

[0135] When this condition is satisfied, data inputted from the ports P2, P3, P4 and P5 is subjected to rate conversion in the buffer memory and changed over in frame unit periodically to thereby be multiplexed by means of the switch, so that the multiplexed data is out-

putted to port P1.

[0136] Accordingly, in order to increase the throughput in writing and reading operation to and from the disk array, it is understood that it is required to be set to be $n \approx m$, that is, the number of ports to be changed over periodically is required to be set to be substantially equal to the ratio of the data transfer rate per channel between the disk adapter and the switch to the data transfer rate per channel between the switch and the disk array.

[0137] For example, a single channel of 4 Gbps is connected between the disk adapter and the switch and four channels of 1 Gbps are connected between the switch and the disk array.

[0138] Further, a single channel of 10 Gbps is connected between the disk adapter and the switch and four channels of 2 Gbps are connected between the switch and the disk array.

[0139] In this case, since the throughput is not balanced between the input and output ports of the switch, the effective throughput is $2\text{Gbps} \times 4 = 8\text{Gbps}$.

[0140] With the above processing, since the rate conversion, the multiplexing and the demultiplexing are performed in the switch SW1, the data transfer rate at the channels D01, D02, D03 and D04 can be increased even if the data transfer rate at the channels D11, D12, D13 and D14 is low.

[0141] That is, the throughput between the disk adapter DKA and the disk array DA can be improved.

[0142] As the data transfer system in the disk storage system of the embodiment, the fibre channel or the InfiniBand can be used.

[0143] Fig. 6 is a diagram illustrating the method of increasing disk drives in the disk storage system of the first embodiment.

[0144] In Fig. 6, disk groups R5 and R6 are increased as compared with Fig. 1.

[0145] In order to increase the disk drives, switches having increased ports are used as the switches SW1 and SW2.

[0146] When the disk drives are increased, the throughput on the disk array side of the switch is increased and the balance of the throughput between the disk adapter side and the disk array side of the switch collapses. Accordingly, there is a possibility that the rate conversion function of the switch is not operated effectively.

[0147] Accordingly, in the switch SW1, a new channel D05 is increased between the disk adapter DKA and the switch SW1.

[0148] Further, in the case of the switch SW2, any new channel is not increased and the signal transmission rate of the channel D02 is increased to thereby balance the throughput between the disk adapter side and the disk array side of the switch.

[0149] For example, in the switch SW1, 8 channels of 1 Gbps are connected between the switch and the disk array and 2 channels of 4 Gbps are connected between the disk adapter and the switch.

[0150] In the switch SW2, 8 channels of 1 Gbps are connected between the switch and the disk array and 1 channel of 10 Gbps is connected between the disk adapter and the switch.

[0151] As described above, in the disk storage system of the embodiment, the disk drives can be increased in accordance with the number of ports of the switch.

[0152] This method of increasing the disk drives can be applied to increase the disk drives of the ATA (AT Attachment) system in which the number of drives capable of being connected to one port is small.

[0153] Fig. 7 schematically illustrates a disk storage system according to a second embodiment of the present invention.

[0154] The disk storage system of the embodiment has the configuration of a disk array portion different from the disk storage system of the first embodiment.

[0155] The disk storage system of the embodiment includes a disk controller DKC and four disk arrays DA1, DA2, DA3 and DA4.

[0156] The disk controller DKC includes a channel adapter CHA, a cache memory CM and a disk adapter DKA.

[0157] The disk array DA1 is connected to the disk adapter DKA through a channel D01 and a switch SW1.

[0158] Similarly, the disk arrays DA2, DA3 and DA4 are connected to the disk adapter DKA through a channel D02 and a switch SW2, a channel D03 and a switch SW3, and a channel D04 and a switch SW4, respectively.

[0159] The switches SW1, SW2, SW3 and SW4 function to perform the rate conversion, the multiplexing and the demultiplexing in the same manner as the first embodiment.

[0160] The fibre channel is used as the data transfer system between the disk adapter DKA and the disk arrays DA1, DA2, DA3 and DA4 through the switches SW1, SW2, SW3 and SW4.

[0161] The switches SW1, SW2, SW3 and SW4 are fibre channel switches.

[0162] The configuration of the disk array in the embodiment is described by taking the disk array DA1 as an example.

[0163] The disk arrays DA1, DA2, DA3 and DA4 have the same drive configuration.

[0164] The disk array DA1 includes a disk array composed of four disks connected to a channel D11, a disk array composed of four disks connected to a channel D12, a disk array composed of four disks connected to the channel D13 and a disk array composed of four disks connected to a channel D14.

[0165] For example, disk drives DK1, DK2, DK3 and DK4 are connected to the channel D11.

[0166] In this manner, as the method of connecting a number of disk drives to one channel to access to a disk drive, there is a fibre channel arbitrated loop (hereinafter referred to FC-AL).

[0167] Fig. 10 shows a connection form of the FC-AL

13

EP 1 353 264 A2

14

by taking a connection form of the disk drives DK1, DK2, DK3 and DK4 as an example.

[0168] Each of input/output ports of the disk drivers and the switch SW1 is provided with a transmitter Tx and a receiver Rx.

[0169] The connection form of the FC-AL is the topology that the input/output ports of the drives and the switch are connected into a loop as shown in Fig. 10, for example.

[0170] The input/output ports of the drives function as NL (Node Loop) ports of the fibre channel.

[0171] The NL port is a port of the apparatus (the disk drive in the embodiment) which makes loop operation.

[0172] The input/output ports of the switch SW1 on the side connected to the disk array DA1 function as FL (Fabric Loop) ports of the fibre channel.

[0173] The FL port is a port of the switch that can connect the FC-AL.

[0174] Since the loop having the FL port functions as a public loop of the fibre channel, the FC-AL formed by the channel D11 is a public loop.

[0175] The public loop is a loop that the disk drive on the loop can communicate with a port outside of the loop through the switch.

[0176] Accordingly, the disk drives DK1, DK2, DK3 and DK4 can communicate with the disk adapter DKA through the switch SW1 and the channel D01.

[0177] As described above, the connection form of the channel D11 has been described by way example, while the channels D12, D13 and D14 are the same as the channel D11.

[0178] In the disk storage system of the embodiment, when the RAID system is constructed, the disk groups R1, R2, R3 and R4 are each the RAID group. In the embodiment, four disk drives constitute the RAID group, although the number of drives constituting the RAID group is not limited to four.

[0179] In the embodiment, the FC-AL is used in the channels D11, D12, D13 and D14 to connect the disk drives.

[0180] Up to 126 disk drives can be connected to each of the channels D11, D12, D13 and D14 in accordance with specification of the FC-AL.

[0181] Further, optical fiber cable is used as a medium of the channels D01, D02, D03 and D04 and metal cable is used as a medium of the channels D11, D12, D13 and D14.

[0182] As described above, in the disk storage system of the embodiment, since the disk drives are connected by means of the FC-AL, the number of connectable drives per port of the switch can be increased.

[0183] In other words, there is the effect that the memory capacity of the disk storage system is increased.

[0184] Further, by connecting the disk drives by means of metal cable, it is not necessary to provide any expensive optical interface in each disk drive and accordingly the cost of the disk drive can be reduced.

[0185] Fig. 8 schematically illustrates a disk storage

system according to a third embodiment of the present invention.

[0186] In the disk storage system of the embodiment, disk controllers and switches are duplicated.

5 [0187] In the embodiment, the fibre channel is used as the data transfer system between disk adapters DKA1, DKA2 and disk array DA1 through switches SW1, SW2.

10 [0188] The disk storage system of the embodiment includes disk controllers DKC1, DKC2, switches SW1, SW2 and disk array DA1.

[0189] The switches SW1 and SW2 function to perform the rate conversion, the multiplexing and the demultiplexing in the same manner as the first embodiment.

15 [0190] The disk controller DKC1 includes channel adapter CHA1, cache memory CM1 and disk adapter DKA1.

[0191] The disk controller DKC2 includes channel adapter CHA2, cache memory CM2 and disk adapter DKA2.

20 [0192] The disk adapter DKA1 is connected to the switch SW1 through channel D1a, the disk adapter DKA2 is connected to the switch SW2 through channel D2a, the disk adapter DKA1 is connected to the switch SW2 through channel D1b and the disk adapter DKA2 is connected to the switch SW1 through channel D2b.

[0193] The disk drives constituting the disk array DA1 each include two input/output ports.

30 [0194] For example, the disk drives DK1, DK2, DK3 and DK4 are connected to both channels D11 and D21.

[0195] The disk array DA1 includes a disk array composed of four disk drives connected to channels D11 and D21, a disk array composed of four disk drives connected to channels D12 and D22, a disk array composed of four disk drives connected to channels D13 and D23 and a disk array composed of four disk drives connected to channels D14 and D24.

35 [0196] The channels D11, D12, D13, D14, D21, D22, D23 and D24 are connected to the disk drives by the FC-AL.

[0197] Fig. 11 shows a connection form of the FC-AL in the embodiment by taking a connection form of the disk drives DK1, DK2, DK3 and DK4 as an example.

40 [0198] Each of the disk drives includes two NL ports.

[0199] A transmitter Tx and a receiver Rx are provided with each of input/output ports of the disk drives and the switches SW1 and SW2.

45 [0200] The input/output ports of the switches SW1 and SW2 on the side connected to the disk array DA1 are FL ports.

[0201] The switch SW1 and the disk drives DK1, DK2, DK3 and DK4 are connected into a loop by means of the channel D11.

50 [0202] Similarly, the switch SW2 and the disk drives DK1, DK2, DK3 and DK4 are connected into a loop by means of the channel D21.

[0203] The two loops are public loops of the fibre

15

EP 1 353 264 A2

16

channel and the disk drives DK1, DK2, DK3 and DK4 can communicate with the disk adapter DKA1 or DKA2 through the switch SW1 or SW2.

[0204] The connection form of the channels D11 and D21 has been described as above by way of example, while the channels D12, D13, D14, D22, D23 and D24 have also the same connection form.

[0205] In the disk storage system of the embodiment, when the RAID system is constructed, the disk groups R1, R2, R3 and R4 are each the RAID group. In the embodiment, four disk drives constitute the RAID group, although the number of drives constituting the RAID group is not limited to four.

[0206] Any of the disk adapters DKA1 and DKA2 can access to all of the disk drives in the disk array DA1.

[0207] The disk storage system of the embodiment uses the channel D1b or D2b as a detour upon failure of the switch SW1 or SW2.

[0208] For example, even when the switch SW1 fails, the disk adapter DKA1 can access to the disk array DA1 through the channel D1b and the switch SW2.

[0209] On the other hand, when the switch SW2 fails, the disk adapter DKA2 can access to the disk array DA1 through the channel D2b and the switch SW1 and accordingly the disk storage system having high reliability can be realized.

[0210] Fig. 12 schematically illustrates a disk storage system according to a fourth embodiment of the present invention.

[0211] In the disk storage system of the embodiment, channels D3a and D3b connecting between the switches SW1 and SW2 are provided in the disk storage system of the third embodiment.

[0212] In the embodiment, the fibre channel is used as the data transfer system between the disk adapters DKA1, DKA2 and the disk array DA1 through the switches SW1, SW2.

[0213] The disk storage system of the embodiment includes disk controllers DKC1, DKC2, switches SW1, SW2 and disk array DA1.

[0214] The switches SW1 and SW2 function to perform the rate conversion, the multiplexing and the demultiplexing in the same manner as the first embodiment.

[0215] The disk controller DKC1 includes channel adapter CHA1, cache memory CM1 and disk adapter DKA1.

[0216] The disk controller DKC2 includes channel adapter CHA2, cache memory CM2 and disk adapter DKA2.

[0217] The disk adapter DKA1 is connected to the switch SW1 through channel D1a, the disk adapter DKA2 is connected to the switch SW2 through channel D2a, the disk adapter DKA1 is connected to the switch SW2 through channel D1b and the disk adapter DKA2 is connected to the switch SW1 through channel D2b.

[0218] Further, the switches SW1 and SW2 are connected to each other through channels D3a and D3b.

[0219] Disk drives constituting the disk array DA1 each include two input/output ports.

[0220] For example, the disk drives DK1, DK2, DK3 and DK4 are connected to both channels D11 and D21.

[0221] The disk array DA1 includes a disk array composed of four disk drives connected to channels D11 and D21, a disk array composed of four disk drives connected to channels D12 and D22, a disk array composed of four disk drives connected to channels D13 and D23 and a disk array composed of four disk drives connected to channels D14 and D24.

[0222] The channels D11, D12, D13, D14, D21, D22, D23 and D24 are connected to the disk drives by the FC-AL as shown in Fig. 11.

[0223] Any of the disk adapters DKA1 and DKA2 can access to all of the disk drives in the disk array DA1.

[0224] In the disk storage system of the embodiment, when the RAID system is constructed, the disk groups R1, R2, R3 and R4 are each the RAID groups. In the embodiment, four disk drives constitute the RAID group, although the number of drives constituting the RAID group is not limited to four.

[0225] The access route between the disk adapters DKA1, DKA2 and the disk array DA1 at the normal state (in case of no failure in switch) is now described.

[0226] The disk adapter DKA1 includes a route (route 1) for accessing to the disk array DA1 through the channel D1a and the switch SW1 and a route (route 2) for accessing to the disk array DA1 through the channel D1b, the switch SW2, the channel D3a and the switch SW1.

[0227] Similarly, the disk adapter DKA2 includes a route (route 3) for accessing to the disk array DA1 through the channel D2a and the switch SW and a route (route 4) for accessing to the disk array DA1 through the channel D2b, the switch SW1, the channel D3b and the switch SW2.

[0228] On the other hand, the channels D1b and D2b are used as detours upon failure of the switch.

[0229] For example, even when the switch SW1 fails, the disk adapter DKA1 can access to the disk array DA1 through the channel D1b and switch SW2.

[0230] On the other hand, when the switch SW2 fails, the disk adapter DKA2 can access to the disk array DA1 through the channel D2b and the switch SW1.

[0231] The throughput between the disk adapter and the disk array in the embodiment is now described.

[0232] As an example, it is assumed that the data transfer rate on the channels D1a, D1b, D2a, D2b, D3a and D3b is 2 Gbps per channel and the data transfer rate on the channels D11, D12, D13, D14, D21, D22, D23 and D24 is 1 Gbps per channel.

[0233] At this time, the total throughput between the switch SW1 and the disk array DA1 is 4 Gbps.

[0234] The total throughput between the disk adapter DKA1 and the switch SW1 is 4 Gbps by using the routes 1 and 2 as the access route.

[0235] Since the throughputs of the switch SW1 on

17

EP 1 353 264 A2

18

the side of disk adapter DKA1 and on the side of disk array DA1 are both 4 Gbps, the throughput between the disk adapter DKA1 and the disk array DA1 is 4 Gbps.

[0236] Similarly, the total throughput between the switch SW2 and the disk array DA1 is 4 Gbps.

[0237] The total throughput between the disk adapter DKA2 and the switch SW2 is 4 Gbps by using the routes 3 and 4 as the access route.

[0238] Since the throughputs of the switch SW2 on the side of disk adapter DKA2 and on the side of disk array DA1 are both 4 Gbps, the throughput between the disk adapter DKA2 and the disk array DA1 is 4 Gbps.

[0239] When the throughput per channel as described above is applied to the third embodiment (Fig. 8), the throughput between the disk adapter DKA1 and the disk array DA1 is limited to the throughput on the channel D1a and is 2 Gbps since the channels D1b and D2b are used as detours only upon failure of the switch.

[0240] Similarly, the throughput between the disk adapter DKA2 and the disk array DA1 is limited to the throughput on the channel D2a and is 2 Gbps.

[0241] In the third embodiment, in order to increase the throughput between the disk adapter and the disk array to 4 Gbps, it is necessary to increase the data transfer rate on both the channels D1a and D2a to 4 Gbps.

[0242] As described above, according to the embodiment, even if the data transfer rate per channel between the disk adapter and the switch is low, the disk storage system having the increased total throughput between the disk adapter and the disk array can be realized.

[0243] As described above, according to the present invention, the following effects can be attained.

[0244] There can be provided the disk storage system having the high throughput between the disk adapter and the disk array.

[0245] Further, there can be provided the disk storage system having the high throughput between the disk adapter and the disk array and the increased disk drives connected in the disk array.

[0246] Moreover, there can be provided the disk storage system having the high reliable disk array.

[0247] Further, there can be provided the disk storage system having the high reliable network between the disk adapter and the disk array.

[0248] Furthermore, there can be provided the disk storage system having the network between the disk adapter and the disk array with high reliability and high throughput.

[0249] Moreover, there can be provided the disk storage system having the high throughput in reading from and writing in the disk.

[0250] Further, there can be provided the disk storage system capable of keeping the high throughput.

[0251] In addition, there can be provided the disk storage system having the high throughput between the disk adapter and the disk array and a low cost.

[0252] It should be further understood by those skilled

in the art that although the foregoing description has been made on embodiments of the invention, the invention is not limited thereto and various changes and modifications may be made without departing from the spirit of the invention and the scope of the appended claims.

Claims

1. A disk storage system including a disk controller (DKC) and a disk array (DA), said disk controller (DKC) including a channel adapter (CHA), a cache memory (CM) and a disk adapter (DKA), wherein:

said disk adapter (DKA) is connected to said disk array (DA) through switches (SW1-4) having buffer memories;

a data transfer rate per channel between said disk adapter (DKA) and said switches (SW1-4) being set to be higher than that between said switches (SW1-4) and said disk array (DA); and said switches (SW1-4) change over connection between ports to which said disk adapter (DKA) is connected and ports to which disk drives constituting said disk array (DA) are connected in accordance with destination information in a frame for each of inputted frames.

2. A disk storage system including a disk controller (DKC) and a plurality of disk arrays (DA1-4), said disk controller (DKC) including a channel adapter (CHA), a cache memory (DA) and a disk adapter (DKA), wherein:

said disk arrays (DA1-4) comprise a plurality of disk drives (DK1-4) connected into a loop; and said disk adapter (DKA) is connected to said plurality of disk arrays (DA1-4) through switches (SW1-4) having buffer memories;

a data transfer rate per channel between said disk adapter (DKA) and said switches (SW1-4) being set to be higher than that between said switches (SW1-4) and said plurality of disk arrays (DA1-4);

said switches (SW1-4) change over connection between ports to which said disk adapter (DKA) is connected and ports to which said plurality of disk arrays (DA1-4) are connected in accordance with destination information in a frame for each of inputted frames.

3. A disk storage system including a disk controller (DKC) and a disk array (DA), said disk controller (DKC) including a channel adapter (DKA), a cache memory (CM) and a disk adapter (DKA), wherein:

said disk adapter (DKA) is connected to said disk array (DA) through switches (SW1-4) hav-

19

EP 1 353 264 A2

20

ing buffer memories; and
 a combination of disk drives connected to the same switch constitutes a RAID group (R1-4);
 a data transfer rate per channel between said disk adapter (DKA) and said switches (SW1-4) being set to be higher than that between said switches (SW1-4) and said disk array (DA);
 said switches change over connection between ports to which said disk adapter (DKA) is connected and ports to which said disk drives constituting said RAID group (R1-4) are connected in accordance with destination information in a frame for each of inputted frames.

4. A disk storage system including first and second disk controllers (DKC1, 2) and a plurality of disk arrays (DA);

said first disk controller (DKC1) including a first channel adapter (CHA1), a first cache memory (CM1) and a first disk adapter (DKA1); and

said second disk controller (DKC2) including a second channel adapter (CHA2), a second cache memory (CM2) and a second disk adapter (DKA2);
 wherein:

said first disk adapter (DKA1) is connected to said plurality of disk arrays (DA) through a first switch (SW1) having a buffer memory and said second disk adapter (DKA2) is connected to said plurality of disk arrays (DA) through a second switch (SW2) having a buffer memory, said first switch (SW1) being connected to said second disk adapter (DKA2), said second switch (SW2) being connected to said first disk adapter (DKA1);

a data transfer rate per channel between said first disk adapter (DKA1) and said first switch (SW1) and between said second disk adapter (DKA2) and said first switch (SW1) being set to be higher than that between said first switch (SW1) and said plurality of disk arrays (DA);

a data transfer rate per channel between said second disk adapter (DKA2) and said second switch (SW2) and between said first disk adapter (DKA1) and said second switch (SW2) being set to be higher than that between said second switch (SW2) and said plurality of disk arrays (DA);

said first switch (SW1) changing over connection between ports to which said first or second disk adapter (DKA1, 2) is connected and ports to which said plurality of disk arrays (DA) are connected in accordance with destination information in a frame for each of inputted frames;
 said second switch (SW2) changing over connection between ports to which said first or second disk adapter (DKA1, 2) is connected and ports to which said plurality of disk arrays (DA)

are connected in accordance with destination information in a frame for each of inputted frames.

5. A disk storage system including first and second disk controllers (DKC1, 2) and a plurality of disk arrays (DA);

said first disk controller (DKC1) including a first channel adapter (CHA1), a first cache memory (CM1) and a first disk adapter (DKA1);

said second disk controller (DKC2) including a second channel adapter (CHA2), a second cache memory (CM2) and a second disk adapter (DKA2);
 wherein:

said first disk adapter (DKA1) is connected to said plurality of disk arrays (DA) through a first switch (SW1) having a buffer memory and said second disk adapter (DKA2) is connected to said plurality of disk arrays (DA) through a second switch (SW2) having a buffer memory, said first switch (SW1) being connected to said second disk adapter (DKA2), said second switch (SW2) being connected to said first disk adapter (DKA1);

a data transfer rate per channel between said first disk adapter (DKA1) and said first switch (SW1) and between said second disk adapter (DKA2) and said first switch (SW1) being set to be higher than that between said first switch (SW1) and said plurality of disk arrays (DA);

a data transfer rate per channel between said second disk adapter (DKA2) and said second switch (SW2) and between said first disk adapter (DKA1) and said second switch (SW2) being set to be higher than that between said second switch (SW2) and said plurality of disk arrays (DA);

said first switch (SW1) being connected to said second switch (SW2) through a channel having the same data transfer rate as that connected between said first disk adapter (DKA1) and said second switch (SW2) and a channel having the same data transfer rate as that connected between said second disk adapter (DKA2) and said first switch (SW1);

said first switch (SW1) changing over connection between ports to which said first or second disk adapter (DKA1, 2) or said second switch (SW2) is connected and ports to which said plurality of disk arrays (DA) are connected in accordance with destination information in a frame for each of inputted frames;

said second switch (SW2) changing over connection between ports to which said first or second disk adapter (DKA1, 2) or said first switch (SW1) is connected and ports to which said plurality of disk arrays (DA) are connected in ac-

21

EP 1 353 264 A2

22

cordance with destination information in a frame for each of inputted frames.

6. A disk storage system according to any of Claims 1 to 5, wherein:

5

data transferred from said disk array (DA) to said switch (SW) is multiplexed in said switch (SW) to be transferred to said disk adapter (DKA) when data is read out from said disk array (DA); and
data transferred from said disk adapter (DKA) to said switch (SW) is demultiplexed in said switch (SW) to be transferred to said disk array (DA) when data is written into said disk array (DA).

10

15

7. A disk storage system according to any of Claims 1 to 5, wherein:

20

said disk adapter (DKA) sets destination information into a frame to be transmitted so that changing over of connection between said ports is made periodically when data is written into said disk array (DA) from said disk adapter (DKA); and
said switch (SW) changes over connection between said ports in accordance with a round robin method when data is read out from said disk array (DA) to said disk adapter (DKA).

25

30

8. A disk storage system according to Claim 7, wherein:

the number of ports to be changed over periodically is set to be substantially equal to a ratio of a data transfer rate per channel between said disk adapter (DKA) and said switch (SW) to that between said switch (SW) and said disk array (DKA).

35

40

9. A disk storage system according to any of Claims 1 to 5, wherein:

said disk adapter (DKA) is connected to said switch (SW) through optical fiber cable and said switch is connected to said disk array through metal cable.


45

50

55

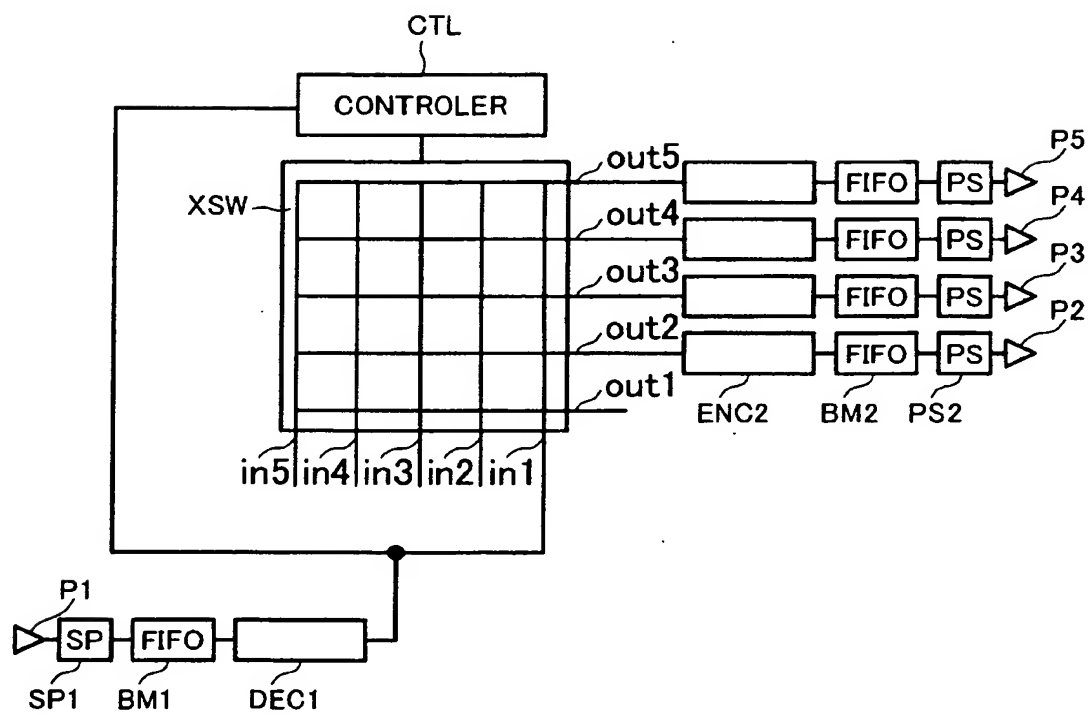
EP 1 353 264 A2

FIG. 1




EP 1 353 264 A2

FIG.2



EP 1 353 264 A2

FIG.3



EP 1 353 264 A2

FIG.4


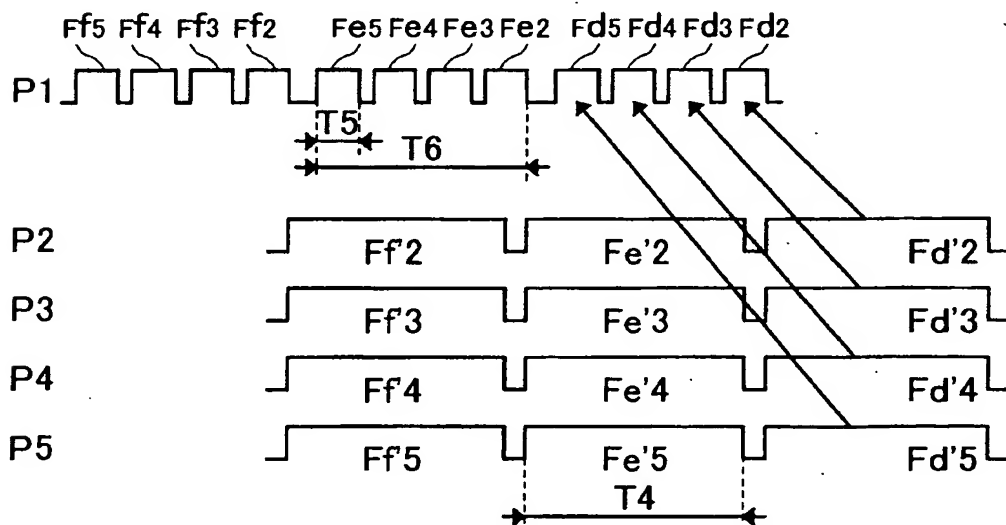


FIG.5



EP 1 353 264 A2

FIG.6





FIG.7



EP 1 353 264 A2

FIG.8






FIG.9




EP 1 353 264 A2

FIG.10




EP 1 353 264 A2

FIG. 11



EP 1 353 264 A2

FIG.12



THIS PAGE BLANK (USPTO)